arXiv:1612.05205v1 [cs.DC] 16 Nov 2016

GentleRain+: Making GentleRain Robust on Clock Anomalies

Mohammad Roohitavaf and Sandeep S. Kulkarni

Department of Computer Science and Engineering
Michigan State University
East Lansing, MI 48824, USA
Email: \{roohitav,sandeep\}@cse.msu.edu

Abstract. Causal consistency is in an intermediate consistency model that can be achieved together with high availability and high performance requirements even in presence of network partitions. There are several proposals in the literature for causally consistent data stores. Thanks to the use of single scalar physical clocks, GentleRain has a throughput higher than other proposals such as COPS or Orbe. However, both of its correctness and performance relay on monotonic synchronized physical clocks. Specifically, if physical clocks go backward its correctness is violated. In addition, GentleRain is sensitive on the clock synchronization, and clock skew may slow write operations in GenlteRain. In this paper, we want to solve this issue in GenlteRain by using Hybrid Logical Clock (HLC) instead of physical clocks. Using HLC, GentleRain protocl is not sensitive on the clock skew anymore. In addition, even if clocks go backward, the correctness of the system is not violated. Furthermore, by HLC, we timestamp versions with a clock very close to the physical clocks. Thus, we can take causally consistency snapshot of the system at any give physical time. We call GentleRain protocol with HLCs GentleRain+. We have implemented GentleRain+ protocol, and have evaluated it experimentally. GentleRain+ provides faster write operations compare to GentleRain that rely solely on physical clocks to achieve causal consistency. We have also shown that using HLC instead of physical clock does not have any overhead. Thus, it makes GentleRain more robust on clock anomalies at no cost.

1 Introduction

Distributed data stores are one of the integral parts of today's Internet services. Service providers usually replicate their data on different nodes worldwide to achieve higher performance and durability. However, when we use this approach, the consistency between replicas becomes a concern. In an ideal situation, all replicas always represent exactly the same data. In other words, any update to any data item instantaneously becomes visible in all replicas. This model of consistency is called *strong consistency*. Strong consistency provides a simple semantics for programmers who want to use the distributed data store for their applications. Unfortunately, it is impossible to achieve strong consistency

without sacrificing the availability when we have network partitions. In particular, in case of network partitions, to maintain the consistency between different copies of data items in different replicas, we have to make updates unavailable. The CAP theorem [12] implies that we can only choose two requirements out of strong consistency, availability, and partition-tolerance.

At the other end, the weakest consistency model is called *eventual consistency* [16]. In this consistency model, as the name suggests, the only guarantee is that replicas become consistent "eventually". We can implement always-available services under this consistency model. However, it may expose clients to anomalies, and application programs need to consider such anomalies when they program their applications. To understand how eventual consistency may leads to anomaly, consider the following example from [15]: Suppose in a social network, Alice changes the status of an album to friends-only. Then, she uploads a personal photo to her album. Under eventual consistency model, the photo may become visible in a remote node before album status. In that case, a user querying a remote node may see Alice's personal photo without being her friend. Despite such anomalies, some distributed data stores use eventual consistency. One example is Dynamo [9] which is used by Amazon to manage the state of some of its services with very high availability requirements.

Causal consistency is an intermediate consistency model. It is weaker than strong consistency, but stronger than eventual consistency. In particular, causal consistency requires that the effect of an event can be visible only when the effect of its causal dependencies is visible. The causal dependency captures the notion of happens-before relation define in [14]. Under this relation, any event by a client depends on all previous events by that client. Thus, in our example, the event of uploading the photo depends on the event of changing the album status to friends-only. Thus, nobody can view Alice's photo before finding her album as friends-only. Causal consistency (with some restrictions as we explain in Section 2) is achievable with availability even in presence of networks partitions.

COPS [15] is one of the first proposals in the community for causally consistent data stores. COPS explicitly maintains the list of dependencies for each version. Then, in a remote replica, we do not make a version visible if some of its dependencies are missing. To check weather all of dependencies of an update are present in the data center, servers communicate with each other. The problem is that the overhead of this explicit dependency checking is high, and it grows as the number of dependencies or the number of serves inside the data center grows. Orbe [10] solves the problem of maintaining an explicit list of dependencies by using dependency matrices, but it still suffers from sending dependency check messages to other partitions.

GentleRain [11] uses a different approach in a sense that it does not send any dependency check message. Instead, it guarantees causal consistent by using physical clocks. In this way, partitions communicate with each other only *periodically*, and they do not need to communicate upon receiving each new update, unlike COPS or Orbe. When compared with COPS and Orbe, this reduces

the message complexity of GentleRain significantly thereby providing higher throughput.

Although GentleRain reduces the overhead of tracking depenendencies, it relies on synchronized and monotonic clocks for both correctness and performance. Specifically, it requires that clocks are strictly increasing. This may be hard to guarantee if the underlying service such as NTP causes non-monotonic updates to POSIX time [7] or suffers from leap seconds [3, 2]. In addition, the clock skew between physical clocks of servers may leads to cases where GentleRain deliberately delays write operations.

This issue is made worse if we have a federated data center. In other words, all the data in the data center is not controlled by the same entity. Such a situation can arise when subsets of data are *controlled* by multiple entities. An example of such federated data center occurs when you have multiple departments in a university that contribute to the data in the data center. In this case, each partition may consist of data of individual department. A client, with appropriate authorization, may be able to access data from multiple/all departments. However, since the partitions are controlled by different entities, enforcing constraints such as tight clock synchronization is difficult. We expect this to be an important issue when the data center is obtained by data from multiple sources and while each source is willing to provide access (via GET/PUT etc) they are not willing to handover the whole data to each other.

Our goal in this paper is making GentleRain more robust on clock anomalies such as clock skew or going backward. To achieve this goal, we focus on the use of hybrid logical clocks [13] that combine the logical clocks [14] and physical clocks. In particular, these clocks assign a timestamp $hlc.e$ to event e such that if e happened before f (as defined in [14]), then $hlc.e < hlc.f$ is true. Furthermore the value of hlc is very close to the physical clock and is backward compatible with NTP clocks [6]. In particular, in [13], it is shown that one can replace the last 16 bits of the NTP clock (that is 64 bits long) by data necessary to model HLC. Moreover, the remaining 48 bits provide precision of microseconds. Thus, an application that uses NTP can rely on HLC without violating its correctness.

Another aspect that we have focused on in this paper is the locality of client requests. Specifically, in GentlRain, clients can only access their local data center. Other systems like COPS also has made that assumption. We provide an impossibility results that proves this assumption is indeed required in any causally consistent data store that makes local updates visible immediately.

Contributions of the paper.

- We show that we can improve the latency of PUT operations in GentleRain with the use of HLC.
- We demonstrate the efficiency provided by our approach by simulating clock skews as well as by performing experiments on cloud services provided by Amazon.
- We demonstrate how one can remove the reliance of GentleRain on clock synchronization service.
- We show using HCL does not have any overhead.

– We provide an impossibility result that shows locality of traffic is necessary for a causally consistent data stores that makes local updates visible immediately.

Organization of the paper. In Section 2, we define our system architecture and the notion of causal consistency. In Section 3, we discuss, in detail, the issues of GentleRain that we want to address. In Section 4, we provide a brief overview of hybrid logical clocks from [13]. Section 5 provides our GentleRain+ and Section 6 provides our experimental results. We provide an impossibility result in Section 7. In Section 8 we discuss why HLC is preferred over physical or logical clocks. Finally, Section 9 conclude the papers.

2 Background

2.1 Architecture

In this section we focus on the system architecture and assumptions that are the same as those assumed in [15], [10], and [11]. We consider data stores that are both partitioned and replicated. Partitioning a data stores means instead of storing all data in a single machine, storing it on several machines. Using partitioning we can have a scalable data store that is not limited by the capacity of a single machine, and can store more data by simply adding more machines. Replication, on the other hand, helps us to increase both performance and durability of our data. Specifically, by replicating data in geographically different locations, we let the clients query their local data center for reading or writing data thereby reducing the response time for client operations. In addition, by replicating data, we increase the durability of our data, as if we lose the data in some replica, we have the data in other data centers. Thus, a data store consists of N partitions each of which is replicated in M replicas.

Figure 1 show a typical architecture of a partitioned and replicated data store. The data is replicated in different data centers in geographically different locations, and inside each data center, data is partitioned over several servers. Note that the partitioning of each replica into partitions is logical. We make no assumptions about physical location of each partition. For example, in an university setting, each partition may consist of data associated with one department. All these partitions form one key-value store. Thus, each partition in the key-value store could be administered independently as long as it provides the necessary functionality of GET and PUT operations. In Figure 1, data center C is a logical data center that uses servers provided by two different departments/providers.

We assume multi-version key-value stores that store several versions for each key. A key-value store has two basic operations: $PUT(k, val)$ and $GET(k)$, where $PUT(k, val)$ writes new version with value val for item with key k, and $GET(k)$ reads the value of a item with key k.

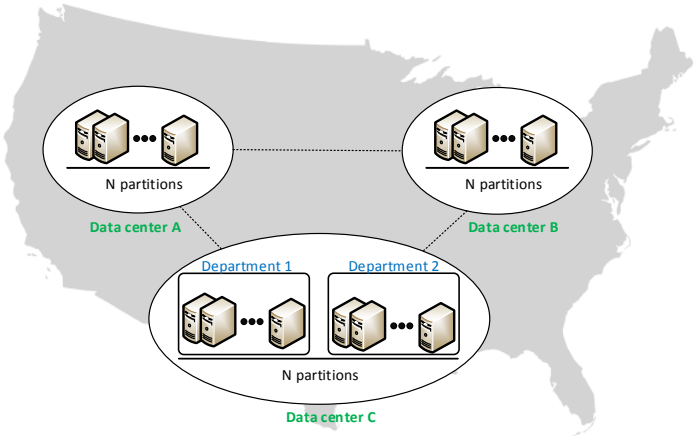

Fig. 1. A typical architecture of a Geo-replicated partitioned data store. Each data center is partitioned into N servers. Data center C is a logical data center consists of servers managed by different departments/providers.

2.2 Causal Consistency

Causal consistency is defined based on the happens-before relation between events [14]. In the context of key-value stores, we define happens-before relation as follows:

Definition 1 (Happens-before). *Let a and b be two events. We say a happens before b, and denote it as $a \rightarrow b$ iff:*

- *a and b are two events by the same client (i.e., in a single thread of execution), and a happens earlier than b, or*
- *a is a $PUT(k, val)$ operation and b is a $GET(k)$ operation and $GET(k)$ returns the value written by a, or*
- *there is another event c such that $a \rightarrow c$ and $c \rightarrow b$.*

Now, we define casual dependency as follows:

Definition 2 (Causal Dependency). *Let v_1 be a version of key k_1, and v_2 be a version of key k_2. We say v_1 causally depends on v_2, and denote it as v_1 dep v_2 iff $PUT(k_2, v_2) \rightarrow PUT(k_1, v_1)$.*

Definition 3 (Visibility). *We say version v of key k is visible to a client c iff $GET(k)$ performed by client c returns v' such that $v = v'$ or $v \rightarrow v'$.*

Now, we define causal consistency as follows:

Definition 4 (Causal Consistency). *Let k_1 and k_2 be any two arbitrary keys in the store. Let v_1 be a version of key k_1, and v_2 be a version of key k_2 such that v_1 dep v_2. The store is causally consistent if whenever v_1 is visible to client c, v_2 is also visible to client c.*

In the above definition we ignore the possibly of conflicts in writes. Conflicts occur when we have two writes on the same key such that there is no causal dependency relation between them. Specifically,

Definition 5 (Conflict). *Let v_1 and v_2 be two versions for key k. We call v_1 and v_2 are conflicting iff $\neg(v_1 \ \mathit{dep} \ v_2)$ and $\neg(v_2 \ \mathit{dep} \ v_1)$. (i.e., none of them depends on the other.)*

In case of conflict, we want a function that resolves the conflict. Thus, we define conflict resolution function as $f(v_1, v_2)$ that returns one of v_1 and v_2 as the winner version. Of course, if v_1 and v_2 are not conflicting, f returns the latest version, i.e., if $v_1 \ \mathsf{dep} \ v_2$ then $f(v_1, v_2) = v_1$. For such function f, Lloyd et al. [15] have named causal consistency with conflict resolution causal+ consistency. We formalize causal+ consistency as follows:

Definition 6 (Causal+ Consistency). *Let k_1 and k_2 be any two arbitrary keys in the store. Let v_1 be a version of key k_1, and v_2 be a version of key k_2 such that $v_1 \ \mathit{dep} \ v_2$. Let f be a conflict resolution function. The store is causal+ consistent for conflict resolution function f if whenever v_1 is visible to client c, then for key k_2, version v_2' is visible to client c as well where $v_2' = v_2$ or $f(v_2, v_2') = v_2'$.*

A trivial solution for a causally (or causal+) consistent data store is a data store that always returns the initial value for each key. Of course, such data store is not desirable. Thus, we define causal++ consistency that requires the data store to make all local updates visible to clients immediately. Specifically,

Definition 7 (Causal++ Consistency). *A store is causal++ consistent for conflict resolution function f if following conditions are hold:*

- *it is causal+ consistent for conflict resolution function f, and*
- *If client c writes version v_1 for key k on replica on r, then version v_2 for key k is immediately visible to any client that accesses replica r where $v_2 = v_1$ or $f(v_1, v_2) = v_2$.*

In practice, in addition to the consistency, we want all replicas eventually converge. In other words, we want a write in a replica to reflect in other connected replicas as well. Thus, we define convergence as follows:

Definition 8 (Convergence). *A store is convergent if the following condition is hold:*

if client c writes version v_1 for key k on replica r, then version v_2 for key k is eventually visible to any client that accesses any replica connected to r where $v_2 = v_1$ or $f(v_1, v_2) = v_2$.

Note that, current proposals for causally consistent data store like [15], [10], and [11] are infact convergent and causal++ consistent. However, they have assumed that clients do not change their replicas. In Section 7, we provide an impossibility result that shows achieving causal++ consistency is impossible if clients access different replicas for the same items.

In the next Section, we focus on our motivation for improving GentleRain.

3 Motivation

Thanks to the use of clocks for dependency checking, GentleRain [11] has a throughput higher than that of other causally consistent data stores proposed in the literature such as COPS [15] or Orbe [10]. In particular, GentleRain avoids sending dependency check messages unlike other systems which leads to a light message complexity which in turn leads to higher throughput. In the following, we review the main parts of GentleRain protocol.

GentleRain. GentleRain assigns each version a timestamp equal to the physical clock of the partition where the write of the version occurs. We denote the timestamp assigned to version X by $X.t$. GentleRain assigns timestamps such that following condition is satisfied:

$C1$: If version X of object x depends on version Y of object y, then $Y.t < X.t$.

Also, each node in the data center periodically computes a variables called Global Stable Time (GST) (through communication with other partitions) such that following condition is satisfied:

$C2$: When GST in a node has a certain value T, then all versions with timestamps smaller than or equal to T are visible in the data center.

When a client preforms $GET(k)$, the partition storing k, returns the newest version v of k which is either created locally, or has a timestamp no greater than GST. According to conditions $C1$ and $C2$ defined above, any version which is a dependency of v is visible in the local data center and causal consistency is satisfied.

Problem: Sensitivity on physical clock and clock synchronization. To satisfy condition $C1$, in some cases, we may have to wait before creating a new version. Specifically, in order to create a new version for a key, the client sends the timestamp t of the last version that it has read/written together with the PUT operation. The partition receiving this request first waits until its physical clock is higher than t, and then creates the new version. The amount of this wait, as we observed in our experiments, is proportional to the clock skew between servers. In other words, as the physical clocks of servers drift from each other, the incidence and the amount of this wait period increases. This makes GentleRain sensitive to clock synchronization between servers, and a drift between clocks of the servers may have a direct negative impact on the response time for the clients operations. This problem becomes more significant when we create a federated data center by aggregating servers managed by different authorities, or even in different locations, as in this case the clock skew between servers is expected to be higher.

In addition to the waiting issue, physical clock must be monotonic for GentleRain to be correct. In other words, if one of the physical clocks goes backward, then the correctness of GentleRain is not guaranteed, and causal consistency guarantee is violated. To understand how that can happen consider this example: consider a system consisting of two data centers A and B. Suppose GSTs

in both data centers are 6. That means, both data centers assume all versions with timestamp smaller than 6 are visible (condition $C2$). Now, suppose the physical clock of one of the servers in data center A goes backward to 5. In this situation, if a client writes a new version at that server, condition $C2$ is violated, as the version with timestamp 5 has not arrived in data center B, but GST is 6 which is higher than 5. Since correctness of GentleRain relays on condition $C2$, by violating $C2$, correctness of GentleRain is violated.

As we explained above, both correctness and performance of GentleRain relay on the accuracy of the physical clocks, and the clock synchronization between servers. We want to eliminate this reliance on physical clocks, and make GentleRain more robust by using HLCs. In the section 4, we overview HLC, and then in Section 5 we focus on using HLC in GentleRain.

4 Hybrid Logical Clocks

In this section, we recall HLCs from [13]. HLC combines logical and physical clocks to leverage key benefits of both. We focus on the benefits of using HLCs in Section 8. A HLC timestamp of event e has two elements: $l.e$ and $c.e$. $l.e$ is the value of the physical clock, and represents our best approximation of the global time when e occurs. And, $c.e$ is a *bounded* counter that is used to capture causality whenever $l.e$ is not enough to capture causality. Specifically, if we have two events e and f such that e happens-before f (see Definition 1), and $l.e = l.f$, to capture causality between e and f, we set $c.e$ to a value higher than $c.f$. Although, we increase c, as it is proved in [13], the theoretical maximum value of c is $O(n)$. In practice, this value remains very small. For completeness, we recall algorithm of HLC from [13] below.

Algorithm 1 HLC algorithm from [13]

1: **Send/Local Event**
2: $l'.a = l.a$
3: $l.a = max(l'.a, pt.a)$ //tracking maximum time event, $pt.a$ is physical time at a
4: **if** $(l.a = l'.a)$ $c.a = c.a + 1$ //tracking causality
5: **else** $c.a = 0$
6: Timestamp event with $l.a, c.a$

7: **Receiving message** m
8: $l'.a = l.a$
9: $l.a = max(l'.a, l.m, pt.a)$ //$l.m$ is l value in the timestamp of the message received
10: **if** $(l.a = l'.a = l.m)$ then $c.a = max(c.a, c.m) + 1$
11: **else if** $(l.a = l'.a)$ then $c.a = c.a + 1$
12: **else if** $(l.a = l.m)$ then $c.a := c.m + 1$
13: **else** $c.a = 0$
14: Timestamp event with $l.a, c.a$

HLC satisfies logical clock property, that it allows us to capture (one-way) causality between two events. Specifically, if e happens-before f, then $HLC.e < HLC.f$, where $HLC.e < HLC.e$ iff $(l.e < l.f \vee ((l.e = l.f) \wedge c.e < c.f)))$. This implies that if $HLC.e = HLC.f$, then e and f are (causally) concurrent. At the same time, just like physical clock, HLC increases spontaneously, and it is close to the physical clock. Thus, it can be used to take snapshot at a given physical time.

5 GentleRain+ Protocol

GentalRain+ is basically GentleRain protocol that uses HLC instead of physical clock. For completeness of algorithms, we provide the algorithms for the basic protocol that supports PUT and GET operation. However, we assume the reader is familiar with GentleRain protocol [11]. Here, we mostly focus on the use of HLC in the protocol. The protocol for transactions (GET-SNAPSHOT and GET-ROTX) is the exactly the same as that of in GentleRain.

5.1 Client Side

The client side of GentleRain+ is identical to the client side of GentleRain, except all timestamps and clock values are HLC values. Table 1 shows variables maintained by a client. A client maintains a dependency time DT_c that is the value of the maximum timestamp of all items accessed by the client. A client also maintains GST_c that is the the maximum GST value that the client is aware of. We recall the client side algorithm from [11] in Algorithm 2.

Table 1. Variables maintained by client c

Variable	Definition
DT_c	dependency time at client c
GST_c	global stable time known by client c

5.2 Server Side

In this section we focus on the server side of the protocol. We have M replicas (i.e., M data centers) each of which with N partitions. We denote the nth partition in mth replica by p_n^m. Table 2 shows variables that define the state of partition p_n^m. PC_n^m and HLC_n^m are the physical clock and the HLC of partition p_n^m, respectively. VV_n^m is vector with size N that maintains the latest (HLC) timestamps received from other replicas by partition p_n^m. Specifically, $VV_n^m[k]$

Algorithm 2 Client operations at client c (These functions are identical to those of [11] except that GST, gst, DT have HLC values instead of physical clock values)

```
 1: GET (key k)
 2:     send ⟨GetReq k, GST_c⟩ to server
 3:     receive ⟨GetReply v, ut, gst⟩
 4:     DT_c ← max(DT_c, ut)
 5:     GST_c ← max(GST_c, gst)
 6: return v

 7: PUT (key k, value v)
 8:     send ⟨PutReq k, v, DT_c⟩ to server
 9:     receive ⟨PutReply ut⟩
10:     DT_c ← max(DT_c, ut)
```

is the latest timestamp received from server p_n^k. LST_n^m is minimum timestamp that p_n^m has received from its peers in other data centers. In other words, LST_n^m maintains the minimum value in VV_n^m. Servers share their LST_n^m with each other, and compute GST as the minimum of LSTs. GST_n^m is the GST computed in server p_n^m.

Table 2. Variables maintained by server p_n^m

Symbol	Definition
PC_n^m	current physical clock time at p_n^m
HLC_n^m	current hybrid logical clock time at p_n^m
VV_n^m	version vector of p_n^m , with M elements
LST_n^m	local stable time at p_n^m
GST_n^m	global stable time at p_n^m

Table 3 shows the information that is stored for each version. For each version, in addition to the key and value, we store some additional metadata including the (HLC) time of creation of the version ut, and the source replica sr where the version has been written.

Algorithm 3 shows the first part of the server side operations of GentleRain+. Upon receiving a GET request (GetReq), the server finds the latest version of the requested key that is either written in the local data center, or its update time is less than the current GST, and then returns back this value together with its update time and GST to the client.

Table 3. Data item

Symbol	Definition
d	item tuple $\langle k, v, ut, , sr \rangle$
k	key
v	value
ut	update time
sr	source replica id

The main difference between GentleRain and GentleRain+ is how we process PUT requests. Upon receiving a PUT request by server p_n^m, the server first updates its HLC by calling function $updateHLCforPut$. This function updates the HLC_n^m by the algorithm of the HLC. We pass dt to $updateHLCforPut$, and $updateHLCforPut$ updates HCL_n^m such that the new HLC_n^m is higher than dt. After updating the HLC_n^m, the server updates its own entry in VV_n^m with the HLC_n^m. Next, the server creates a new version for the item specified by the client and uses the current HLC_n^m value for its timestamp. The server, next, sends back the assigned timestamp $d.ut$ to the client. Note that in GentleRain, processing PUT operation is delayed deliberately if dt sent by the client is higher than the current physical clock. Such delay is not needed in GentleRain+.

Upon creating new version for an item in one data center, we should replicate the new version to other replicas in other data centers. Thus, we call $sendReplicate$ function that sends replicate messages to other data centers. Upon receiving a replicate message from server p_n^k, the receiving server p_n^m adds the new item to the version chain of key $d.k$. The server also updates the entry for server p_n^k in its version vector. Thus, it sets $VV_n^m[k]$ to $d.ut$.

Algorithm 4 shows the second part of the GentleRain+ protocol. The functions in Algorithm 4 (except $updateHLC$) are the same as those of GentleRain except the use of HLC instead of physical clock. Heartbeat messages are sent by a server, if the server has not sent any replicate message for a certain time Δ. The goal of heartbeat messages is updating the knowledge of the peers of an partition in other replicas (i.e., updating VVs). In addition, every θ time, we compute LST and GST. To compute GST, partitions needs to communicate their LST with each other. Broadcasting LSTs has a high overhead. Instead, we efficiently compute GST over a tree as describe in [11].

6 Experimental Results

We have implemented GenlteRain+ protocol in a distributed kay-value store called MSU-DB. MSU-DB is written in Java, and it can be downloaded from [4]. MSU-DB uses Berkeley DB [8] in each server for data storage and retrieval. For the server-to-server communication it relays on Netty [5]. For comparison purposes, we have implemented GentleRian in the same code base.

First, we investigate the overhead of using HLC. Next, we focus on the response time improvement for PUT operations resulted by using HLC instead of

Algorithm 3 Server operations at server p_n^m (part 1)

Upon receive $\langle$GetReq $k\rangle$
  $GST_n^m \leftarrow max(GST_n^m, gst)$
  obtain latest version d from version chain of key k s.t.

- $d.sr = m$, or
- $d.ut \leq GST_n^m$

  send $\langle$GetReply $d.v, d.ut, GST_n^m\rangle$ to client

Upon receive $\langle$PutReq $k, v, dt\rangle$
  $updateHCLforPut(dt)$
  $VV_n^m[m] \leftarrow HLC_n^m$
  Create new item d
  $d.k \leftarrow k$
  $d.v \leftarrow v$
  $d.ut \leftarrow HLC_n^m$
  $d.sr \leftarrow m$
  insert d to version chain of k
  send $\langle PutReply\ d.ut\rangle$ to client
  sendReplicates(d)

Upon receive $\langle$Replicate $d\rangle$ from p_n^k
  insert d to version chain of key $d.k$
  $VV_n^m[k] \leftarrow d.ut$

sendReplicates(d)
  for each server $p_n^k, k \in \{0 \dots M-1\}, k \neq m$ **do**
    send $\langle$Replicate $d\rangle$ to p_n^k

updateHLCforPut (dt)
  $l = HLC_n^m.l$
  $HLC_n^m.l = max(l', PC_n^m, dt.l)$
  if $(HLC_n^m.l = l' = dt.l)$ $HLC_n^m.c \leftarrow max(HLC_n^m.c, dt.c) + 1$
  else if $(HLC_n^m.l = l')$ $HLC_n^m.c \leftarrow HLC_n^m.c + 1$
  else if $(HLC_n^m.l = l)$ $HLC_n^m.c \leftarrow dt.c + 1$
  else $HLC_n^m.c = 0$

Algorithm 4 Server operations at server p_n^m (part 2) (These functions are identical to those of [11] except we use HLC instead of physical clock)

Upon every Δ time
  if more than Δ time unite has passed since the last heartbeat or replicate message
    $updateHCL()$
    $VV_n^m[m] \leftarrow HLC_n^m$
    for each server $p_n^k, k \in \{0 \dots M-1\}, k \neq m$ **do**
      send $\langle$HEARTBEAT $HLC_n^m\rangle$ to p_n^k

Upon receive $\langle$HEARTBEAT $hlc\rangle$ from p_n^k
  $VV_n^m[k] \leftarrow hlc$

Upon every θ time
  $LST_n^m \leftarrow min_{i=1}^{M}(VV_n^m[i])$
  $GST_n^m \leftarrow min_{j=1}^{N}(LST_j^m)$

updateHLC ()
  $l = HLC_n^m.l$
  $HLC_n^m.l = max(HLC_n^m.l, PC_n^m)$
  if $(HLC_n^m.l = l)$ $HLC_n^m.c \leftarrow HLC_n^m.c + 1$
  else $HLC_n^m.c \leftarrow 0$

physical clocks. We run our experiments on Amazon AWS [1] on `c3.large` instances running Ubuntu 14.04. The specification of servers is as follows: 7 ECUs, 2 vCPUs, 2.8 GHz, Intel Xeon E5-2680v2, 3.75 GiB memory, 2 x 16 GiB Storage Capacity.

6.1 Overhead of HLC

In this section, we investigate the overhead of using HLC instead of physical clocks. We use the compact representation of HLC suggested in [13]. Thus, we do not need to use two different variables for two parts of a HLC clock. Specifically, we replace the last 16 bits of the NTP clock (that is 64 bits long) by data necessary to model HLC. The remaining 48 bits provide precision of microseconds.

Using compact representation, HLC timestamps are that same as physical clock timestamps regarding storage and memory requirements. The only concern is the processing overhead of updating HLC timestamps for PUT operations. Although HLC update is a very simple operation, to make sure that it does not affect the throughput of the system we perform the following experiment: we run MSU-DB on a single server, and run enough clients to saturate the server with PUT operations. For each of these PUT operations, the server needs to update the HLC. We run the experiment for PUT operations with different value sizes and compute the throughput as the number of operation done in one second.

We do the same experiment on GentleRain. Note that, running the experiment on a single server is in favor of GentleRain, as the wait period before PUT operation does not occur at all, and GentleRain can respond PUT operations immediately. We also implement the echo protocol in the same code base that simply echos the client messages. We compute the throughput of the echo server to investigate the processing power of our hardware and software platform. Table 4 shows the results of our experiment. We observed that using HLC instead of physical clocks does not have any overhead on the throughput.

Table 4. Throughput in Kop/sec for PUT operations with different value sizes.

	Echo (16)	**PUT (16)**	**PUT (128)**	**PUT (1K)**
GentleRain	51.612066	10.053134	7.12	1.5823334
GentleRain+	51.612066	10.4336	7.0791333	1.5866

6.2 Response Time of PUT Operations

In this section we study the effect of clock skew on the response time for PUT operations in GentleRain. We also investigate how HLC eliminates the negative effect of clock skew on the response time in GentleRain+.

To study the effect of clock skew on the response time accurately, we need to have a precise clock skew between servers. However, the clock skew between two different machines depends on many factors out of our control. To have a more accurate experiment, we consolidate two virtual servers on a single machine, and impose an artificial clock skew between them. Then, we change the value of the clock skew and observe its effect of the response time for PUT operations. A client sends PUT requests to the servers in a round robin fashion. Since the physical clock of one of the servers is behind the physical clock of the other server, half of the PUT operations will be delayed by the GentleRain. On the other hand, GentleRain+ does not delay any PUT operation, and processes them immediately. We compute the average response time for 1K PUT operations. Figure 2 shows average response time for PUT operation in GentleRain grows linearly as the clock skew grows, while the average response time in GentleRain+ is independent of clock skew.

Next, we do the same the experiment when the servers are running on two different machines. This time, we do not impose any artificial clock skew. We run NTP [6] on servers to synchronized physical clocks. Now, the client sends PUT requests to these servers. We repeat this experiment when the serves are running in different availability zones in the Virginia region. We always run servers in the same availability zone, and the client in a different availability zone in the same region. Table 5 shows average delay before PUT operations in GentleRain. We do not have such delay in GentleRain+. Figure 3 shows how this delays negatively

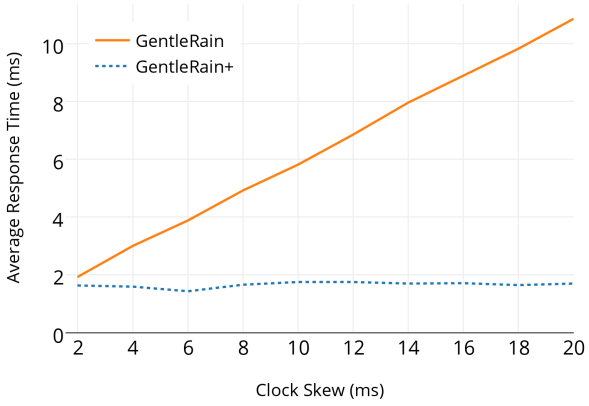

Fig. 2. The effect of clock skew on the average response time for PUT operations.

affect the average response time of GentleRain. The increase of average response time directly affect the clients.

Table 5. Average delay before PUT operations in GentleRain.

Availability Zone	Average Delay (ms)	Ping from client to servers (ms)
us-east-1b	7.75	1.261
us-east-1c	12.25	1.973
us-east-1e	12.5	0.794

7 Impossibility Results

In this section, we want to focus on the locality of traffic assumption made by GentleRain. The CAP theorem [12] proves that achieving strong consistency, together with availability is impossible under the asynchronous network model. Availability means any client operation (PUT and GET) should be answered in a finite time. Asynchronous network model means that the communication among servers can be arbitrary delayed and the network can be partitioned for an unbounded duration.

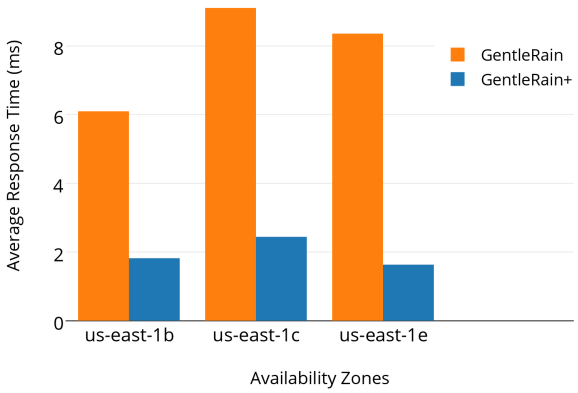

Fig. 3. The comparison of average response time between GentleRain and GentleRain+ in different availability zones.

Although achieving strong consistency together with availability is impossible in asynchronous network model, there are several proposals in the literature to guarantee causal++ consistency together with availability under asynchronous network model. However, they assume the locality of traffic. In other words, they assume that a client accesses the same replica forever. In this section, we consider the *moving client model* where client can access different keys from different replicas.

Next, we claim that the locality of traffic assumption is essential for causal++ consistency, as we have the following impossibility result:

Theorem 1. *In an asynchronous network with moving client model, it is impossible to implement a replicated data store that guarantees following properties:*

- *Availability*
- *Causal++ consistency*

Proof. We prove this by contradiction. Assume an algorithm A exists that guarantees availability and causal++ consistency for conflict resolution function f in asynchronous network with moving client model. We create an execution of A that contradicts causal++ consistency. Assume a data store where each key is stored in at least two replicas r and r'. Initially, the data center contains two keys k_1 and k_2 with versions v_1^0 and v_2^0. These versions are replicated on r and r'. Next the system executes in the following fashion.

- There is a partition between r and r', i.e., all future messages between them will be lost.

- c performs $PUT(k_1, v_1^1)$ on replica r.
- c performs $PUT(k_2, v_2^1)$ on replica r.
- c' performs $GET(k_2)$ on replica r. Value returned is v_2^1
- c' moves to replica r'.
- c' performed $GET(k_1)$ on replica r'. Value returned is v_1^0.

Thus, v_2^1 is visible to c', but its causally dependent version v_1^1 is not visible. One the other hand, there is no v_3 on replica r' such that $f(v_1^1, v_3) = v_3$. Thus, causal++ consistency is violated (contradiction).

8 Discussion: Why Hybrid Logical Clocks

The main property of logical clock proposed by Lamprot [14] is that we can use logical clock to capture causality between events. Specifically, if event e happens-before f (see Definition 1), then logical clock timestamp assigned to e is smaller than logical clock timestamp assigned to f. Although logical clocks help us to capture causality between events, using them in the distributed data store system has some complications. First, as mentioned in [11], using logical clocks can cause clocks at different servers move at very different rates. Specifically, the logical clock of a server with higher event rate moves faster than the logical clock of a server with smaller event rate. When clocks of different servers moves at different speed, the visibility of update can be delayed arbitrary [11]. Another problem of logical clocks is the timestamp assigned to version has not meaningful relation with a *real clock*. Thus, if we use logical clocks for timestamping versions, we have no way to access a certain version of a key or snapshot of the system at the certain physical time.

Physical clocks, on the other hand, do not have the above mentioned problem of logical locks. Specifically, they advance spontaneously. Thus, unlike logical clocks, even if no event occurs at a server, its clock advances by its own. Another advantage of physical clock is that if use them to timestamp versions, we can access the data of our system at a desired physical time. Although physical clocks have these benefits, they are not as efficient as logical clocks for capturing causal relation between events. In particular, if we timestamp events with physical clocks as soon as they occur, then it is not guaranteed that if event e happens-before f timestamp of e is smaller than timestamp of f. To solve this problem, the approach of GentelRain is delaying write operations, but it can leads to increasing response time when clock synchronization is not good enough. Thus, it makes it sensitive on clock synchronization.

Hybrid logical clocks combine the benefit of both logical clocks and physical clocks. Specifically, at one hand, just like physical clocks, it advances spontaneously, and it also is close to the real clock. On the other hand, without delaying event, it guarantees logical clock property that is if event e happens-before f timestamp of e is smaller than timestamp of f.

9 Conclusion

In the paper we provided an improvement over GentleRain protocol to make it more robust on clock anomalies. The correctness of GentleRain relays on monotonic clocks. Thus, if =clocks go backward, the correctness of GentleRain is violated. In addition, GentleRain may delay write operation, because of clock skew between servers. When clocks are synchronized with each other this delay does not occur. However, if for any reason clocks become unsynchronized, the forced delay adversely affect the response time of write operations. Thus, both correctness and performance of GentleRain relay on the synchronized monotonic clocks. Our improvement over GentleRain, name GentelRain+, solves both problems by using HLC instead of physical clock. Thus, if clocks go backward the correctness is still satisfied. In addition, clock drift does not affect the performance of the system at all, as all operations are processed immediately unlike GentleRain.

We implemented GentelRain+ protocol in a distributed key-value store called MSU-DB, and compared it with GentleRain. As expected, GentleRain+ provides shorter response time for write operations comparing to GentleRain. We also did experiments to investigate the overhead of using HLC instead of physical clock. We observed using HLC does not have any processing overhead comparing to physical clocks. In addition, as we used compact representation of HLC, the memory and storage overhead of HLC is exactly the same as physical clock. Thus, GentleRain+ make GentleRain more robust at no cost.

Finally, we provided an impossibility result which states in presence of network partitions the locality of traffic is necessary, to have an always available causally consistent data store that immediately makes local updates visible.

References

1. Amazon aws. `https://aws.amazon.com/`.
2. Another round of leapocalypse. `http://www.itworld.com/security/288302/another-round-leapocalypse`.
3. The future of leap seconds. `http://www.ucolick.org/œsla/leapsecs/onlinebib.html`.
4. Msu-db. `http://cse.msu.edu/~roohitav/msudb`.
5. Netty. `http://netty.io/`.
6. The network time protocol. `http://www.ntp.org/`.
7. The trouble with timestamps. `http://aphyr.com/posts/299-the-trouble-with-timestamps`.
8. Berkley db, 2016.
9. Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. Dynamo: Amazon's highly available key-value store. *SIGOPS Oper. Syst. Rev.*, 41(6):205–220, October 2007.
10. Jiaqing Du, Sameh Elnikety, Amitabha Roy, and Willy Zwaenepoel. Orbe: Scalable causal consistency using dependency matrices and physical clocks. In *Proceedings of the 4th Annual Symposium on Cloud Computing*, SOCC '13, pages 11:1–11:14, New York, NY, USA, 2013.

11. Jiaqing Du, Călin Iorgulescu, Amitabha Roy, and Willy Zwaenepoel. Gentlerain: Cheap and scalable causal consistency with physical clocks. In *Proceedings of the ACM Symposium on Cloud Computing*, SOCC '14, pages 4:1–4:13, New York, NY, USA, 2014.
12. Seth Gilbert and Nancy Lynch. Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59, June 2002.
13. Sandeep S. Kulkarni, Murat Demirbas, Deepak Madappa, Bharadwaj Avva, and Marcelo Leone. *Logical Physical Clocks*. 2014.
14. Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. *Commun. ACM*, 21(7):558–565, July 1978.
15. Wyatt Lloyd, Michael J. Freedman, Michael Kaminsky, and David G. Andersen. Don't settle for eventual: Scalable causal consistency for wide-area storage with cops. In *Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, SOSP '11, pages 401–416, New York, NY, USA, 2011.
16. Werner Vogels. Eventually consistent. *Commun. ACM*, 52(1):40–44, January 2009.